

INJECTIVE MAPS BETWEEN FLIP GRAPHS

JAVIER ARAMAYONA, THOMAS KOBERDA, AND HUGO PARLIER

ABSTRACT. We prove that every injective simplicial map $\mathcal{F}(S) \rightarrow \mathcal{F}(S')$ between flip graphs is induced by a subsurface inclusion $S \rightarrow S'$, except in finitely many cases. This extends a result of Korkmaz–Papadopoulos which asserts that every automorphism of the flip graph of a surface without boundary is induced by a surface homeomorphism.

1. INTRODUCTION

Consider a compact, connected and orientable surface S , of genus $g \geq 0$ with $b \geq 0$ boundary components. Moreover, assume that S has $p + q > 0$ marked points, with $p \geq 0$ in the interior of S and the other $q \geq 0$ in ∂S , subject to the condition that every component of ∂S must contain at least one marked point. When convenient, we sometimes think of marked points in the interior of S as punctures.

By an *arc* on S we will mean the homotopy class (relative to the marked points) of an arc properly contained in S , and which intersects the set of marked points only at its endpoints. A *multiarc* is a collection of arcs on S with pairwise disjoint interiors. A maximal multiarc is called a *triangulation*; observe that every triangulation of S contains exactly $d(S) = 6g + 3b + 3p + q - 6$ arcs. The *flip graph* $\mathcal{F}(S)$ is the simplicial graph whose vertices are triangulations of S , and where two triangulations are adjacent if and only if they share exactly $d(S) - 1$ arcs; note this implies that the remaining two arcs intersect exactly once. The reader may wish to note the equivalence of the flip graph and the *triangulation graph* used by other authors, such as M. Bell [2].

Observe that $\mathcal{F}(S)$ is locally finite, as every vertex has valence at most $d(S)$. Since the mapping class group $\text{Mod}(S)$ acts on $\mathcal{F}(S)$ by automorphisms and since the quotient is compact, it follows by the Švarc–Milnor Lemma (see, e.g., [3]) that $\mathcal{F}(S)$ and $\text{Mod}(S)$ are quasi-isometric. This fact has been exploited by Disarlo–Parlier [6] to give an elementary proof of a result of Masur–Minsky [9] that subsurface inclusions induce quasi-isometric embeddings between the corresponding mapping class groups. In a different direction, the flip graph has recently been used by Costantino–Martelli [4] to construct families of quantum representations of mapping class groups.

In this paper, we classify all injective simplicial maps between flip graphs. Before giving a precise statement we need some definitions. Given surfaces

Date: December 1, 2014.

S and S' , by an *embedding* of S into S' we mean a π_1 -injective continuous map $h : S \rightarrow S'$ that maps every marked point on S to a marked point on S' .

An embedding $h : S \rightarrow S'$ induces an injective simplicial map $\phi : \mathcal{F}(S) \rightarrow \mathcal{F}(S')$ as follows: we choose a triangulation A of $S' \setminus \text{int}(h(S))$, plus a collection B of arcs on $\partial h(S)$ whose union is homeomorphic to $\partial h(S)$, and then define $\phi(v) = h(v) \cup A \cup B$ for all $v \in \mathcal{F}(S)$.

The purpose of this paper is to prove that, provided S is “complicated enough”, every injective simplicial map $\mathcal{F}(S) \rightarrow \mathcal{F}(S')$ arises in this way. More concretely, we will say that the surface S is *exceptional* if it is an essential subsurface of (and possibly equal to) a torus with at most two marked points, or a sphere with at most four marked points. Our main result is:

Theorem 1.1. *Suppose S is non-exceptional. Then every injective simplicial map $\phi : \mathcal{F}(S) \rightarrow \mathcal{F}(S')$ is induced by an embedding $S \rightarrow S'$.*

Note that if S is a cylinder with two boundary components and one vertex on each boundary then $\mathcal{F}(S) \cong \mathbb{R}$ with its usual simplicial structure, and thus the statement of Theorem 1.1 is false in this case. If S is a torus with one marked point then $\mathcal{F}(S)$ is an infinite trivalent tree, which we conjecture can be embedded in the flip graph of any surface of genus ≥ 2 with one marked point in a way that is not induced by an embedding between the corresponding surfaces. While these examples highlight the failure of Theorem 1.1 for an arbitrary surface S , we do not know at this time whether Theorem 1.1 holds for some of the surfaces excluded in the hypotheses.

Theorem 1.1 should be compared with a previous result [1] of the first author, which shows the analogous statement for injective maps between pants graphs of surfaces. While the proofs of both results are similar in spirit, the technicalities are rather different; that said, we suspect that it should be possible to give axiomatic conditions for certain classes of graphs, built from arcs or curves on surfaces, that guarantee that simplicial injections between two such graphs are always induced by embeddings of the underlying surfaces.

As usual, the flip graph becomes a geodesic metric space by declaring the length of each edge to be equal to 1. Combining Theorem 1.1 with the convexity result of Disarlo–Parlier [6], we obtain the following corollary:

Corollary 1.2. *Suppose S is non-exceptional, and let $\phi : \mathcal{F}(S) \rightarrow \mathcal{F}(S')$ be an injective simplicial map. Then $\phi(\mathcal{F}(S))$ is a totally geodesic subset of $\mathcal{F}(S')$; in other words, any geodesic in $\mathcal{F}(S')$ connecting two points of $\phi(\mathcal{F}(S))$ is entirely contained in $\phi(\mathcal{F}(S))$.*

We now give an idea of the proof of Theorem 1.1. The first step will be to show the following:

Theorem 1.3. *Suppose S is non-exceptional, and let $\phi : \mathcal{F}(S) \rightarrow \mathcal{F}(S')$ be an injective simplicial map. Then:*

- (1) We have $d(S) \leq d(S')$;
- (2) There exists a multiarc $A \subset S'$, with $d(S') - d(S)$ elements, such that $A \subset \phi(v)$ for all $v \in \mathcal{F}(S)$.

In other words, $\phi(\mathcal{F}(S)) \subset \mathcal{F}_A(S')$, where $\mathcal{F}_A(S')$ denotes the subgraph of $\mathcal{F}(S')$ spanned by those triangulations of S' that contain A . Observe that there is a natural isomorphism $\mathcal{F}_A(S') \cong \mathcal{F}(S' \setminus A)$, where $S' \setminus A$ is the result of cutting S' open along every element of A , and thus we can view the map ϕ as a simplicial injection $\mathcal{F}(S) \rightarrow \mathcal{F}(S' \setminus A)$. Noting that $d(S) = d(S' \setminus A)$, Theorem 1.1 will follow from:

Theorem 1.4. *Let S and S' be connected surfaces, with $d(S) = d(S')$ and S non-exceptional. Then every injective simplicial map $\mathcal{F}(S) \rightarrow \mathcal{F}(S')$ is induced by a homeomorphism $S \rightarrow S'$.*

Remark. Theorem 1.1 was previously shown by Korkmaz-Papadopoulos [8] in the special case when $S = S'$, $\partial S = \emptyset$, and the map $\mathcal{F}(S) \rightarrow \mathcal{F}(S')$ an automorphism.

Acknowledgements. The authors would like to thank the Institute for Mathematical Sciences of Singapore, the Mathematisches Forschungsinstitut Oberwolfach, the Technion, and the Universidad de Zaragoza, where parts of this work were completed. We would like to thank the referee for his/her comments.

The first named author was supported by BQR and Campus Iberus grants. The second named author was partially supported by NSF grant DMS-1203964. The third author is supported by Swiss National Science Foundation grant PP00P2-128557.

2. PATHS IN $\mathcal{F}(S)$

Similarly to the case of pants graphs [1], a large part of our arguments boil down to understanding when it is possible to extend a pair of adjacent edges in $\mathcal{F}(S)$ to a *square* or a *pentagon*; see below for definitions. As it turns out, this issue is significantly more subtle here than for pants graphs, due to the fact that there are vertices in $\mathcal{F}(S)$ with non-isomorphic links. The purpose of this section is to prove a series of technical results that will overcome these difficulties. We begin with some definitions.

2.1. Flippable vs. unflippable arcs. Let $v \in \mathcal{F}(S)$ be a triangulation, and let $a \subset S$ be an arc such that $a \in v$. We will say that a is *flippable with respect to v* if there exists a triangulation $v' \in \mathcal{F}(S)$ that is adjacent to v in $\mathcal{F}(S)$ and satisfies $v \setminus (v \cap v') = a$. In other words, the edges in the triangulations v and v' of S differ only by the arc a . We will denote the flip from v to v' by $a \rightarrow a'$. Observe that $a \in v$ is unflippable if and only if, up to a homeomorphism of S , v contains the arcs in Figure 1; furthermore, if v' is any other triangulation containing those arcs, then a is unflippable with respect to v' .




FIGURE 1. An unflippable arc of a triangulation

Given a vertex $v \in \mathcal{F}(S)$, denote by $\deg(v)$ the *valence* of v in $\mathcal{F}(S)$; that is, the number of vertices of $\mathcal{F}(S)$ that are adjacent to v . Observe that $\deg(v) \leq d(S)$ for every vertex $v \in \mathcal{F}(S)$, and that there exists a vertex $u \in \mathcal{F}(S)$ for which $\deg(u) = d(S)$; indeed, in the light of the previous paragraph it suffices to consider a triangulation that contains no arcs bounding a once-punctured disk.

2.2. Squares and pentagons. A *square* (resp. a *pentagon*) in $\mathcal{F}(S)$ is a closed path with four (resp. five) vertices. Korkmaz-Papadopoulos [8] have shown that every square and pentagon in $\mathcal{F}(S)$ is of the form of the one described in Figures 2 and 3; see Lemmas 2.2 and 2.3 of [8]. In particular we have the following observation, which we state as a separate lemma.

Lemma 2.1. *Let σ be a square or pentagon in $\mathcal{F}(S)$. Then $\bigcap_{v \in \sigma} v$ consists of exactly $d(S) - 2$ curves.*




FIGURE 2. Schematics of a square in the flip graph

As mentioned above, one of the main difficulties stems from the fact that there exist pairs of adjacent edges of $\mathcal{F}(S)$ that are not contained in a square or a pentagon. In light of this, we introduce the notion of an *extendable* edge: we will say that the (oriented) edge (u, v) of $\mathcal{F}(S)$ is extendable if for all $w \in \text{link}(v) \setminus \{u\}$, there is a square or a pentagon in $\mathcal{F}(S)$ that contains $\{u, v, w\}$; here, $\text{link}(v)$ denotes the link of the vertex v in $\mathcal{F}(S)$, i.e. the set of vertices adjacent to v . A path is called *extendable* if it consists of extendable edges.

The next two propositions will be crucial for the proof of Theorem 1.3, as they guarantee that there are enough extendable paths in $\mathcal{F}(S)$.




FIGURE 3. Schematics of a pentagon in the flip graph

Proposition 2.2. *Every pair of vertices $u, v \in \mathcal{F}(S)$ with*

$$\deg(u) = \deg(v) = d(S)$$

may be joined by a path $u = u_0, u_1, \dots, u_n = v$ such that (u_i, u_{i+1}) is extendable for all $0 \leq i < n$.

Proposition 2.3. *Let $u, v \in \mathcal{F}(S)$ be adjacent vertices with $\deg(u) > \deg(v)$. For every $w \in \text{link}(v) \setminus \{u\}$ either u, v, w belong to a common square or there is an extendable path between u and w .*

The proof of the above propositions is somewhat involved and will be broken down into a series of intermediate lemmas. We need some notation before commencing. By a *triangle* on S we mean a simply connected region bounded by three arcs on S . We define *quadrilaterals*, *pentagons* and *hexagons* in a similar fashion. Finally, by a *cylinder* we mean an essential subsurface of S that is homeomorphic to $\mathbb{S}^1 \times [0, 1]$, without any interior marked points, and whose boundary components are two distinct arcs on S .

Lemma 2.4. *Let (u, v) and (v, w) be adjacent edges in $\mathcal{F}(S)$ with*

$$\deg(u) \geq \deg(v).$$

Then either (u, v) and (v, w) are contained in a common square or a pentagon, or the two corresponding flips from u to v and from v to w are supported inside a common cylinder.

Proof. Denote by $a \rightarrow a'$ and $b \rightarrow b'$ the flips from u to v and from v to w , respectively, and observe that $b \neq a'$. If a' and b belong to two different triangles of v then u, v, w are contained in a square in $\mathcal{F}(S)$, which has $v' := (u \setminus b) \cup b'$ as remaining vertex.

Suppose now that a' and b belong to the same triangle of v . In this case b is contained in a quadrilateral containing a and a' . We denote the remaining

sides of this square by c, d, e as in Figure 4. Note that $\deg(u) \geq \deg(v)$ implies that the arc a does not bound a once-punctured disk; otherwise an unflippable arc in u would become flippable in v . In turn, this yields that $c \neq d$ and $b \neq e$.

If $b \neq c$ then there is a second triangle of v to which b belongs. The union of this triangle with the former quadrilateral is a pentagon, and hence one can extend the edges (u, v) and (v, w) to a pentagon in $\mathcal{F}(S)$ as in Figure 5. If, on the other hand, $b = c$, then we see that u, v, w are connected by two flips that are supported in a common cylinder. This completes the proof of the lemma. \square




FIGURE 4. A quadrilateral in S




FIGURE 5. Completing edges to a pentagon in $\mathcal{F}(S)$

A flip that is supported on a cylinder will be called a *cylinder flip*, see Figure 6. The following lemma states that it is always possible to “bypass” a cylinder flip:

Lemma 2.5. *Let S be a non-exceptional surface and let $u, v \in \mathcal{F}(S)$ be adjacent vertices such that $\deg(u) = \deg(v)$ and the flip from u to v is a cylinder flip. Then there exists a path $u = u_1, u_2, \dots, u_k = v$ with*

$$\deg(u_i) = \deg(u) = \deg(v)$$

and such that the flip from u_i to u_{i+1} is a non-cylinder flip for all $i \in \{1, \dots, k-1\}$.

Remark. If $u, v \in \mathcal{F}(S)$ are adjacent vertices of equal degree and the flip from u to v is not a cylinder flip, then such flip is supported on a quadrilateral on S whose boundary arcs are pairwise distinct; compare with Figures 1 and 6.

Proof of Lemma 2.5. Let u and v be adjacent vertices with $\deg(u) = \deg(v)$ and such that the flip from u to v is a cylinder flip. Then u and v contain the arcs the left and right pictures of Figure 6, respectively; we will refer to the labeling shown therein.




FIGURE 6. Local schematics for a cylinder flip.

Observe that if $a = b$ or if both a and b are contained in ∂S , then S is exceptional. We can thus assume that $a \not\subset \partial S$ and $a \neq b$. In particular a belongs to two distinct triangles on u and we have the following pentagon as in Figure 7.




FIGURE 7. A pentagon in S , with α belonging to two distinct triangles

If $c = d$ then $b \not\subset \partial S$, for otherwise S is exceptional. We can thus consider the other triangle to which b belongs, and denote its other edges by e and f . Again, if $e = f$, then S is exactly a four times punctured sphere and is

thus exceptional. It follows that $e \neq f$. As the setup is symmetric, this in turn implies that $c \neq d$. We therefore have three possibilities:

Case 1: $c, d \neq b$. In this case, the original cylinder flip from u to v can be completed to a pentagon by considering the flips inside the pentagon bounded by c, d, b and the two copies of l (see Figure 8). The path we need between u and v is the path inside this pentagon that connects u to v in a clockwise manner. By the remark before the proof, the only remaining issue is to check that all the flips performed along this path are made on quadrilateral edges which are distinct in S . This may be verified, for instance, by looking at Figure 8.




FIGURE 8. The pentagon in $\mathcal{F}(S)$ when $c, d \neq b$

Case 2: $d = b$. As cannot be a torus with a single boundary component and a single marked point on the boundary, we have $c \notin \partial S$. Therefore c belongs to a second triangle; we denote the two remaining edges of this triangle by c' and c'' , as in Figure 9. Again, as S is not exceptional (and in this case is not a torus with two marked points), we have that $c' \neq c''$. We now have a hexagon formed by the different edges of the triangles, and we flip in this hexagon following the schematics in Figure 10.





FIGURE 9. A hexagon in S , where $d = b$



 FIGURE 10. Flipping inside a hexagon to get a path in $\mathcal{F}(S)$

As before, it suffices to check that each flip is performed in a quadrilateral with four distinct edges in S , and we claim that this is the case. The hexagon in S under consideration has edges labeled cyclically by $\{d, l, d, l, c'', c'\}$. We include the diagonal being flipped under the cylinder flip and obtain a sequence of flips as is illustrated by Figure 10. The cylinder flip we wish to avoid is the top flip in the clockwise direction. The sequence of flips we use to avoid the cylinder flip traverses the heptagon in the counterclockwise direction. The claim follows immediately by inspection.

Case 3: $c = b$. We argue similarly – this time $d \not\subset \partial S$, so we consider the other triangle to which d belongs. As before, because S is not exceptional, the two other sides of this triangle are necessarily distinct. Flipping in the resulting hexagon, we find a path (see Figure 11) in which it is easy to check again that all flips are performed in quadrilaterals with distinct sides. This finishes the proof. \square

We can now give proofs of the propositions.

Proof of Proposition 2.2. In light of Lemmas 2.4 and 2.5, it suffices to show that any two vertices of maximal degree can be connected by a path γ , all of whose vertices have maximal degree. So take two vertices u and v of degree


FIGURE 11. Flipping in another hexagon to get a path in $\mathcal{F}(S)$

$d(S)$ and a path

$$\gamma_0 : u = u_1, \dots, u_n = v$$

of minimal length between them. We will proceed by induction on n .

Suppose, for contradiction, that γ_0 contains a vertex of non-maximal degree, and let $1 \leq i \leq n$ be the smallest index such that $\deg(u_i) < d(S)$. The flip from u_{i-1} to u_i has local schematics as illustrated in Figure 12. We denote the loop surrounding the central vertex by a .

FIGURE 12. Encountering a non-maximal degree vertex on γ_0

As the degree of v is maximal, the loop a in Figure 12 must be flipped again along γ_0 . If the arc resulting from flipping a remains surrounded by the two arcs b and c in Figure 12, then the result is locally described by the left side of Figure 12. However, the minimality of n implies such a sequence of flips cannot occur, for otherwise γ_0 could be shortened to a path with

the same endpoints and length $n - 2$. In particular, at least one edge of γ_0 between u_i and v must flip one of the edges $\{b, c\}$. The local behavior of the resulting triangulation u' obtained by flipping one of $\{b, c\}$ is given by Figure 13.




FIGURE 13. The local behavior of the triangulation u'

Observe that by construction, $d_{\mathcal{F}(S)}(v, u') < d_{\mathcal{F}(S)}(v, u_i)$, where $d_{\mathcal{F}(S)}$ denotes the combinatorial distance in $\mathcal{F}(S)$. Now, Figure 14 gives a path to u' which avoids introducing the arc a until the last flip.




FIGURE 14. Modifying γ_0 to avoid a before u'

The previous two paragraphs furnish a path γ_1 between u and v with the following two properties: u' is a vertex of γ_1 , and no triangulation corresponding to a vertex on γ_1 between u and u' contains the arc a .

We repeat the construction above, starting with the flip corresponding to the edge of γ_1 which resulted in the triangulation u' . By induction on n , after $k \leq n$ modifications, we will obtain a path γ_k from u to v , each vertex of which other than the penultimate (i.e. the vertex immediately preceding v) does not contain the arc a . From Figure 12, it follows that we may delete the last two edges of γ_k to obtain a path from u to v in which every triangulation avoids a .

Because there are only finitely many possible unflippable arcs occurring along γ_0 , and because none of the modifications performed to produce γ_k from γ_0 introduce any new unflippable arcs, we may perform finitely many further modifications of γ_k in order to obtain a path γ from u to v in which each triangulation has no unflippable arcs. \square

Proof of Proposition 2.3. The proposition will follow from an argument similar to that of Proposition 2.2. Locally, the flip from u to v is modeled on Figure 12.

Suppose first $w \in \text{link}(v) \setminus \{u\}$ is obtained by flipping along an arc a which is not portrayed in Figure 12; in particular, $w = (v \setminus a) \cup a'$. Then $v' := (u \setminus a) \cup a'$, w , u and v form a square in $\mathcal{F}(S)$ and we are done.

Suppose now that $w \in \text{link}(v) \setminus \{u\}$ is obtained by flipping along an arc which is portrayed in Figure 12. By symmetry we can suppose that it is the lower arc b in such a configuration. It follows that the triangulation w locally looks as in Figure 13.

Now by Lemma 2.5 and Proposition 2.2, it suffices to show that we can connect u and w by a path whose every vertex has degree $\deg(u)$. This is exactly one of the steps in the previous proof — in fact, the path we need is the one shown in Figure 14. \square

Remark. In fact, the arguments above show that, for every “realizable” valence $d \leq d(S)$, any two vertices of valence d in $\mathcal{F}(S)$ may be joined by a path whose every vertex has valence at least d and whose every edge is extendable.

3. PROOF OF THEOREM 1.3

Proof of part (1) of Theorem 1.3. Choose a vertex $u \in \mathcal{F}(S)$ with $\deg(u) = d(S)$. Since ϕ is injective, it follows that

$$d(S) = \deg(u) \leq \deg(\phi(u)) \leq d(S'),$$

the desired conclusion. \square

3.1. The invariant multiarc. The rest of the section is devoted to proving part (2) of Theorem 1.3. Let $u \in \mathcal{F}(S)$ be a vertex with $\deg(u) = d(S)$, and let $u_1, \dots, u_{d(S)}$ be its neighbors. Define

$$A(u) = \phi(u) \cap \phi(u_1) \cap \dots \cap \phi(u_{d(S)});$$

that is, $A(u)$ is the collection of those arcs of u that are not flipped when passing from u to u_i , with $i = 1, \dots, d(S)$. Observe that, by construction, $A(u) \subset \phi(v)$ for every $v \in \text{link}(u)$. Since ϕ is injective, we have that $\phi(u_i) \neq \phi(u_j)$ if $i \neq j$, and thus $A(u)$ consists of exactly $d(S') - d(S)$ vertices.

Our first step is to prove the following:

Lemma 3.1. *Let $u, v \in \mathcal{F}(S)$ be vertices with $\deg(u) = \deg(v) = d(S)$. Then $A(u) = A(v)$.*

Proof. In light of Propositions 2.2 and 2.3, it suffices to prove the result in the case where u and v are adjacent and where (u, v) is an extendable edge. Suppose this is the case, noting that $A(u) \subset \phi(v)$ by the construction of $A(u)$. It suffices to prove that $A(u) \subset \phi(w)$ for every w adjacent to v . As the edge (u, v) is extendable, then u, v, w are contained in a square or pentagon in $\mathcal{F}(S)$, which we denote τ . Let z be the vertex of τ adjacent to

u and not equal to v , noting that $A(u) \subset \phi(z)$, since u and z are adjacent. Now, the injectivity of ϕ implies that $\phi(\tau)$ is a square or a pentagon in $\mathcal{F}(S')$. It follows that the vertices of $\phi(\tau)$ have exactly $d(S') - 2$ arcs in common, by Lemma 2.1. Now,

$$A(u) \subset \phi(z) \cap \phi(u) \cap \phi(v) = \phi(z) \cap \phi(u) \cap \phi(v) \cap \phi(w),$$

so that $A(u) \subset \phi(w)$ as well, as we set out to prove. \square

Denote the multiarc associated to some (and hence any, by Lemma 3.1) vertex of $\mathcal{F}(S)$ of valence $d(S)$ by A . As an immediate corollary of the proof of Lemma 3.1, we have:

Corollary 3.2. *Let $u \in \mathcal{F}(S)$ be a vertex such that $A \subset \phi(v)$ for every $v \in \text{link}(u)$. If the edge (u, v) is extendable, then $A \subset \phi(w)$ for every $w \in \text{link}(v)$.*

We are finally in a position to prove part (2) of Theorem 1.3.

Proof of part (2) of Theorem 1.3. Let v be a vertex of $\mathcal{F}(S)$; we want to prove that $A \subset \phi(v)$. Choose a vertex u of valence $d(S)$. By Propositions 2.2 and 2.3, there exists a path

$$u = u_0, u_1, \dots, u_n = v$$

with the property that the edge (u_i, u_{i+1}) is extendable for all $0 \leq i < n$. Note that $A \subset \phi(w)$ for every $w \in \text{link}(u)$, by Lemma 3.1 and the construction of A . Applying Corollary 3.2 and induction, we obtain that $A \subset \phi(u_i)$ for all $0 \leq i \leq n$. This finishes the proof of Theorem 3.1. \square

4. PROOF OF THEOREM 1.4

We now give a proof of Theorem 1.4. Let S and S' be connected surfaces, with S non-exceptional and $d(S) = d(S') = d$. Let

$$\phi : \mathcal{F}(S) \rightarrow \mathcal{F}(S')$$

be an injective simplicial map. The proof consists of several steps:

4.1. Inducing a map on arcs. We first explain how the map ϕ induces a map

$$\psi : \mathcal{A}(S) \rightarrow \mathcal{A}(S')$$

between the corresponding arc graphs. Here, the arc graph of a surface Z is the simplicial graph $\mathcal{A}(Z)$ whose vertices are arcs on Z , and where two arcs are adjacent in $\mathcal{A}(Z)$ if and only if they have disjoint interiors.

Let $a \in \mathcal{A}(S)$ and consider $\mathcal{F}_a(S)$. Since $\mathcal{F}_a(S) \cong \mathcal{F}(S \setminus a)$ and $d(S \setminus a) = d - 1$, Theorem 1.3 applied to

$$\phi : \mathcal{F}_a(S) \rightarrow \mathcal{F}(S')$$

implies that there exists a unique arc b on S' such that $\phi(\mathcal{F}_a) \subset \mathcal{F}_b$. We set

$$\psi(a) := b.$$

We state the next observation as a lemma, as we will need to make use of it later. In what follows, $i(\cdot, \cdot)$ denotes geometric intersection number between arcs.

Lemma 4.1. *The maps ϕ and ψ satisfy the following properties:*

- (1) *Let $v = (a_1, \dots, a_d) \in \mathcal{F}(S)$. Then $\phi(v) = (\psi(a_1), \dots, \psi(a_d))$.*
- (2) *Let $a \neq a' \in A(S)$ with $i(a, a') = 0$. Then $\psi(a) \neq \psi(a')$ and $i(\psi(a), \psi(a')) = 0$.*
- (3) *Let $a, a' \in A(S)$. If $i(a, a') = 1$ then $i(\psi(a), \psi(a')) = 1$.*

Proof. Part (1) is an immediate consequence of the construction of ψ . For part (2), one may first extend a and a' to a triangulation of S and then apply (1). Finally, to see (3) note that there exist adjacent vertices $v, v' \in \mathcal{F}(S)$ such that $a \in v$ and $a' \in v'$. Since $\phi(v)$ and $\phi(v')$ are also adjacent, the result now follows from (1). \square

In the light of part (1) of Lemma 4.1, Theorem 1.4 will follow once we prove:

Proposition 4.2. *The map $\psi : \mathcal{A}(S) \rightarrow \mathcal{A}(S')$ is an isomorphism.*

Indeed, once Proposition 4.2 has been established, Theorem 1.4 will follow as a combination of Lemma 4.1 (1) and Theorem A of Disarlo [5], which states that two arc graphs are isomorphic if and only if the underlying surfaces are homeomorphic; we remark that this result is originally due to Irmak-McCarthy [7] in the particular case when $\partial S = \emptyset$ and $S' = S$.

4.2. Proof of Proposition 4.2. The key ingredient in the proof of Proposition 4.2 is the following:

Lemma 4.3. *Let $v \in \mathcal{F}(S)$. If $a \in v$ is unflippable, then $\psi(a) \in \phi(v)$ is unflippable.*

Before proving Lemma 4.3, observe that if two arcs a, a' are edges of the same triangle of a triangulation v , then $\psi(a)$ and $\psi(a')$ are edges of the same triangle on S' . To see this, note that two arcs a, a' are edges of the same triangle if and only if there exists an arc $b \in S$ such that $i(a, b) = i(a', b) = 1$ and $i(b, c) = 0$ for all $c \in v \setminus \{a, a'\}$. The desired conclusion now follows from parts (2) and (3) of Lemma 4.1.

In particular, ψ induces a map from the set triangles determined by v to the set of triangles determined by $\phi(v)$. Moreover, if two triangles share a given arc $a \in \mathcal{A}(S)$, then their images under ψ share the arc $\psi(a)$.

Proof of Lemma 4.3. Let v be a triangulation, and suppose $a \in v$ is unflippable, so v contains the arcs a and b depicted in Figure 1, as we draw again in Figure 15. Extend a and b to a triangulation v' that contains the solid arcs in Figure 15. Consider the triangulation v'' obtained from v' by flipping b to the arc b' of Figure 15; as such, v'' determines two triangles Δ and Δ' that share exactly two arcs, namely a and b' .

By the preceding paragraphs and by Lemma 4.1, the triangles $\psi(\Delta)$ and $\psi(\Delta')$ share exactly two arcs, namely $\psi(a)$ and $\psi(b')$. Lemma 4.1 implies that ψ preserves the property of having intersection number 0 (resp. 1), so that we may conclude that $i(\psi(b), \psi(b')) = 1$ and that $i(\psi(b), \psi(c)) = 0$ for every arc $c \in \Delta \cup \Delta'$ with $c \neq b'$. In particular, $\psi(b)$ bounds a once punctured disk on S' whose interior contains the interior of $\psi(a)$; in other words, $\psi(a)$ is unflippable with respect to $\phi(v')$, and hence is unflippable with respect to $\phi(v)$. \square




FIGURE 15. An unflippable arc, revisited

Observe that a vertex $v \in \mathcal{F}(S)$ has maximal valence if and only if all arcs of v are flippable with respect to v . As a consequence of Lemma 4.3, we obtain:

Corollary 4.4. *For all $v \in \mathcal{F}(S)$, $\deg(\phi(v)) = \deg(v)$. In particular, ϕ is surjective and is thus an isomorphism.*

Proof. For the first part, observe that Lemma 4.3 implies $\deg(\phi(v)) \leq \deg(v)$; the other inequality is immediate, as ϕ is injective.

To prove that ϕ is surjective, let $w \in \mathcal{F}(S')$. By the connectivity of $\mathcal{F}(S)$, we may assume that w is adjacent to $\phi(v)$, for some $v \in \mathcal{F}(S)$. Since $\deg(\phi(v)) = \deg(v)$, we have that $w = \phi(v')$ for some v' adjacent to v , as desired. \square

We are finally in a position to prove Theorem 1.4, which as we noted follows from Proposition 4.2:

Proof of Proposition 4.2. First, since ϕ is surjective, part (1) of Lemma 4.1 implies that ψ is surjective as well. We claim that ψ is also injective. Let $a, b \in \mathcal{A}(S)$ with $a \neq b$. If $i(a, b) = 0$ then $\psi(a)$ and $\psi(b)$ are distinct and disjoint, again by Lemma 4.1 (1), and hence we are done. Thus assume $i(a, b) \neq 0$; equivalently, $\mathcal{F}_a(S) \cap \mathcal{F}_b(S) \neq \emptyset$. Using the fact that ϕ is an isomorphism, we obtain

$$\emptyset = \phi(\mathcal{F}_a(S) \cap \mathcal{F}_b(S)) = \phi(\mathcal{F}_a(S)) \cap \phi(\mathcal{F}_b(S)) = \mathcal{F}_{\psi(a)}(S') \cap \mathcal{F}_{\psi(b)}(S'),$$

and thus $i(\psi(a), \psi(b)) \neq 0$, which is what we set out to prove. \square

5. PROOF OF THEOREM 1.1

We are now in a position to prove our main result. Suppose S is a non-exceptional surface, and let $\phi : \mathcal{F}(S) \rightarrow \mathcal{F}(S')$ be an injective simplicial map. By Theorem 1.3, there exists a multiarc A on S' , with $d(S') - d(S)$ elements, such that $A \subset \phi(v)$ for every vertex v of $\mathcal{F}(S)$. In other words, $\phi(\mathcal{F}(S)) \subset \mathcal{F}_A(S')$; recall that $\mathcal{F}_A(S')$ denotes the subgraph of $\mathcal{F}(S')$ spanned by those triangulations of S' that contain A .

Now, $\mathcal{F}_A(S') \cong \mathcal{F}(S' \setminus A)$, where $S' \setminus A$ is the surface obtained from S' by cutting open along every element of A ; note that $d(S) = d(S' \setminus A)$. The surface $S' \setminus A$ need not be connected; let $\Sigma_1, \dots, \Sigma_n$ be its connected components. By slight abuse of notation, the map ϕ induces a map

$$\phi : \mathcal{F}(S) \rightarrow \mathcal{F}(\Sigma_1) \times \dots \times \mathcal{F}(\Sigma_n).$$

Write $\pi_i : \mathcal{F}(\Sigma_1) \times \dots \times \mathcal{F}(\Sigma_n) \rightarrow \mathcal{F}(\Sigma_i)$ for the projection onto the i -th factor.

Claim. *Up to reordering of the indices, $\pi_i \circ \phi$ is trivial for all $i = 2, \dots, n$.*

Proof of the claim. Using the same notation as in the previous section, let

$$\psi : \mathcal{A}(S) \rightarrow \mathcal{A}(\Sigma_1) \times \dots \times \mathcal{A}(\Sigma_n)$$

be the map on arcs graphs induced by ϕ , and choose $a \in \mathcal{A}(S)$. Up to reordering indices, we may assume that $\psi(a) \subset \Sigma_1$. We claim that if $b \in \mathcal{A}(S)$ satisfies $i(a, b) = 0$, then $\psi(b) \subset \Sigma_1$ as well. Indeed, it suffices choose a third arc c such that $i(a, c) = i(b, c) = 1$, so that $i(\psi(a), \psi(c)) = i(\psi(b), \psi(c)) = 1$ by Lemma 4.1. The claim now follows from since $\mathcal{A}(S)$ is connected. \square

In the light of the claim above, we may view ϕ (again abusing notation) as a map

$$\phi : \mathcal{F}(S) \rightarrow \mathcal{F}(\Sigma_1),$$

noting that $d(S) = d(\Sigma_1)$. Since S is not exceptional and connected, it follows that ϕ is induced by a homeomorphism $S \rightarrow \Sigma_1$ by Theorem 1.4. This finishes the proof of Theorem 1.1.

REFERENCES

- [1] J. Aramayona. Simplicial embeddings between pants graphs. *Geom. Dedicata* 144 (2010), 115–128.
- [2] M. Bell. An algorithm for deciding reducibility. Preprint.
- [3] M. R. Bridson, A. Haefliger. *Metric spaces of nonpositive curvature*. Grundlehren der Mathematischen Wissenschaften, 319. Springer-Verlag, Berlin, 1999. xxii+643 pp.
- [4] F. Costantino, B. Martelli. An analytic family of representations for the mapping class group of punctured surfaces. *Geom. Topol.* 18 (2014), no. 3, 1485–1538.
- [5] V. Disarlo. Combinatorial rigidity of arc complexes. Preprint.
- [6] V. Disarlo, H. Parlier. The geometry of flip graphs and mapping class groups. Preprint.
- [7] E. Irmak, J. McCarthy, Injective simplicial maps of the arc complex, *Turkish Journal of Mathematics*, 34 (2010), no. 3, 339–354.

- [8] M. Korkmaz, A. Papadopoulos. On the ideal triangulation graph of a punctured surface. *Ann. Inst. Fourier (Grenoble)* 62 (2012), no. 4, 1367–1382.
- [9] H. Masur, Y. Minsky. Geometry of the complex of curves II. Hierarchical structure. *Geom. Funct. Anal.* 10 (2000), no. 4, 902–974.

INSTITUT DE MATHÉMATIQUES DE TOULOUSE, UNIVERSITÉ PAUL SABATIER, 118 ROUTE
DE NARBONNE, 31062 TOULOUSE, CEDEX 4, FRANCE
E-mail address: `Javier.Aramayona@math.univ-toulouse.fr`

DEPARTMENT OF MATHEMATICS, YALE UNIVERSITY, 20 HILLHOUSE AVE, NEW HAVEN,
CT 06520, USA
E-mail address: `thomas.koberda@gmail.com`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF FRIBOURG, CHEMIN DU MUSÉE 23,
CH-1700 FRIBOURG, SWITZERLAND
E-mail address: `hugo.parlier@unifr.ch`